The Importance of “Spillovers” in the Policy Mission of the Advanced Technology Program*

Adam B. Jaffe
Brandeis University and National Bureau of Economic Research
Waltham, MA 02254

Abstract

Government policies like the Advanced Technology Program (“ATP”) are intended, at least in part, to remedy the “market failure” inherent in the fact that a significant portion of the social benefits of new knowledge and technology are not captured by a firm that invests in R&D. ATP’s project selection, and its evaluation of the impact of its program, can be made more effective by explicitly incorporating the analysis of such “spillovers.” For project selection, this means identifying technological, organizational and economic factors that tend to point to a large “spillover gap,” or deviation between the social and private rates of return to a proposed project. For program evaluation and assessment, it means adapting existing study methods that measure social returns to innovation in ways that explicitly capture spillover effects.

Economic Analysis of R&D Spillovers

Sources of Spillovers

Spillovers have been of interest to economists since at least the nineteenth century. Alfred Marshall, one of the founders of modern microeconomics, argued that R&D spillovers were on the rise, remarking, “the secrecy of business is on the whole diminishing, and the most important improvements in method seldom remain secret for long after they have passed from the experimental stage” (Marshall 1920).

Analytically, it is useful to distinguish several different mechanisms by which R&D generates spillovers. For convenience, I refer to these as “knowledge spillovers,” “market spillovers,” and “network spillovers.” In order to think through the implications of spillovers for ATP, it is useful to consider each of these separately, and then to note that they also interact in a way that tends to increase their combined effect.

Knowledge Spillovers. The quote from Marshall refers to the phenomenon of knowledge spillovers. Knowledge created by one agent can be used by another without compensation, or with compensation less than the value of the knowledge. Knowledge spillovers are particularly likely to result from basic research, but they are also produced by applied research and technology development. This can occur in obvious ways, such as “reverse engineering” of products, and also in less obvious ways, such as when one firm’s abandonment of a particular research line signals to others that the line is unproductive and hence saves them the expense of learning this themselves. The spillover beneficiary may use the new knowledge to copy or imitate the commercial products or processes of the innovator, or may use the knowledge as an input to a research process leading to other new technologies.

In some circumstances the creation of knowledge spillovers is intentional on the part of the innovator; the publication of scientific papers is, at least in part, intended to spread new knowledge so that it can be used by the widest possible audience. In the case of patented inventions, society requires disclosure of new knowledge as a quid pro quo for the granting of monopoly rights in the commercial use of an invention. The effect of this disclosure is, in principle, to make the new knowledge available to others for the purpose of facilitating new and different applications, while at the same time protecting the inventor against copying.

More generally, commercial development and use of new knowledge will tend to cause it to spread, despite any desire of the inventor to prevent such spread. Economic exploitation of new knowledge requires the sale of new products or the incorporation of new processes into commercial use. Such commercialization tends, in general, to reveal at least some aspects of the new knowledge to other economic agents. Hence the very process of economically exploiting the knowledge that research creates tends to pass that knowledge to others. Because the spread of knowledge is greatly affected by the commercial use of new

* This paper is based on a study that I performed for the ATP, Economic Analysis of Research Spillovers: Implications for the Advanced Technology Program, NIST GCR 97-708. I have benefited from comments and useful discussions with Zvi Griliches, Jeanne Powell, Rosalie Ruegg, and Richard Spivack. Some of the ideas in this paper grew out of previous joint research with James Adams. The views expressed herein are my own, however, and should not be attributed to any of these individuals or to the ATP.

Journal of Technology Transfer Vol. 23 (2): 11–19.

technology, even the analysis of "knowledge" spillovers must be informed by an understanding of the market mechanisms that govern the spread of new technology.

Market Spillovers. Market spillovers result when the operation of the market for a new product or process causes some of the benefits thereby created to flow to market participants other than the innovating firm. It is this "leakage" of benefits through the operation of market forces, rather than the flow of knowledge itself, that distinguishes market spillovers from knowledge spillovers. Any time a firm creates a new product, or reduces the cost of producing an existing product, the natural operation of market forces will tend to cause some of the benefits thereby created to be passed on to buyers.

Consider first the case of new or improved products. It is likely that a firm that sells a better mousetrap will charge a price that is higher than that being charged for ordinary mousetraps. But innovative products, even those that are patented or otherwise protected from direct competition, will generally be sold at prices that do not fully capture all of the superiority of the new product relative to what was available before. As a result, consumers will be made better off by the introduction of the new product. This increase in consumer welfare is a social benefit from a new product that is not captured by the innovator. Similarly, if a company does R&D to lower its production cost, it will typically lower its selling price as a result. Again, the innovator's customers are better off, and a benefit is created that is not captured by the innovator. Of course, innovation often results in both higher quality and lower prices; thereby benefiting customers even more.

Network Spillovers. Network spillovers result when the commercial or economic value of a new technology is strongly dependent on the development of a set of related technologies. An example of network spillovers exists among all of the different developers of application software for use with a new operating system platform. If one firm develops a particular application, people will buy it only if many other firms develop other sufficient applications so that the platform itself is attractive and widely used. The term "network spillover" is chosen because the different related research projects are like the different users of a network. The value of a network to any one participant is an increasing function of the number of participants; here the expected value of any one research project is an increasing function of the number of different projects undertaken.

If the commercial payoff to each of a set of related research projects is dependent on all or a significant fraction of the projects being completed successfully, then private firms might hesitate to undertake any one of the projects, for fear that the others will not be undertaken. Conversely, if any one firm decides to undertake such a project, it creates a positive externality for all the other firms, by increasing the probability that the "critical mass" will be achieved. Note that this positive externality or spillover exists even if there is no knowledge spillover among the firms (although it is likely that knowledge spillovers would also be occurring).

The existence of network externalities creates a "coordination problem" that is another possible market failure associated with research. Where network externalities are important, it is possible that firms' inability to coordinate their efforts will lead to a misdirection of research effort, away from the activities associated with network externalities, even if firms are in the aggregate undertaking a socially efficient level of research effort.

It is important to emphasize that the coordination problem only arises if there are reasons why a single firm cannot develop all of the necessary related components (or contract with others for their development) and thereby internalize the network externality. Thus while you cannot run a computer without an operating system, the need for the operating system software does not create an important coordination problem, as the hardware manufacturer can either write the operating system itself or contract for its creation. What distinguishes the operating system (which does not create a significant network externality) from the need for applications programs (which might) is the likelihood that many different applications will ultimately be necessary, and that it is unlikely that one firm would have all the capabilities to create all of these different applications, or even to know what the set of necessary applications will ultimately look like. To put it differently, synergistic market interactions among a small number of technologies is unlikely to create a coordination problem, but when the number of technologies that must be developed is large and the necessary capabilities are diverse, the coordination problem may become severe.

There are a number of different mechanisms by which the coordination problem created by network externalities can be handled. Research joint ventures, in which a number of companies combine forces, can be used to pursue the interrelated approaches whose commercial success is interdependent. By fostering the creation of such joint ventures, ATP assists this process. In addition, the formation of focused programs targeted at a set of interrelated technologies can be used to try to ensure that a critical mass will be reached. Discussion of focused programs is beyond the scope of this paper, but is discussed in Jaffe (1996).

Private and Social Returns to R&D

Pure Market Spillovers. As noted above, the effect of spillovers is to create a gap between the private rate of return to R&D (the return or profit earned by the firm undertaking the research) and the social rate of return, which includes the private return but also includes benefits to the firms' customers and to other firms. The nature of this spillover gap in the context of market and knowledge spillovers is illustrated in Figures 1, 2, and 3.

Figure 1 illustrates a "pure" market spillover. If "Firm

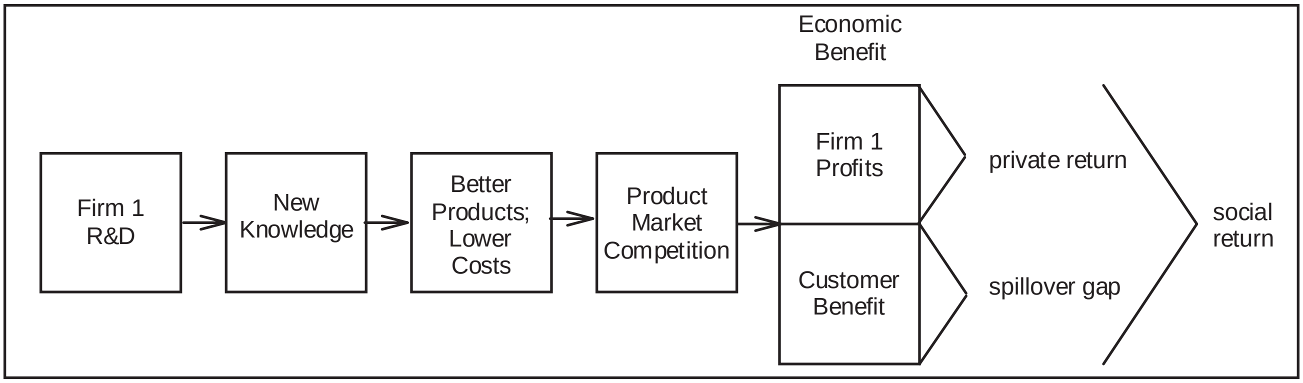

Figure 1. Private and social returns to R&D: Pure market spillover

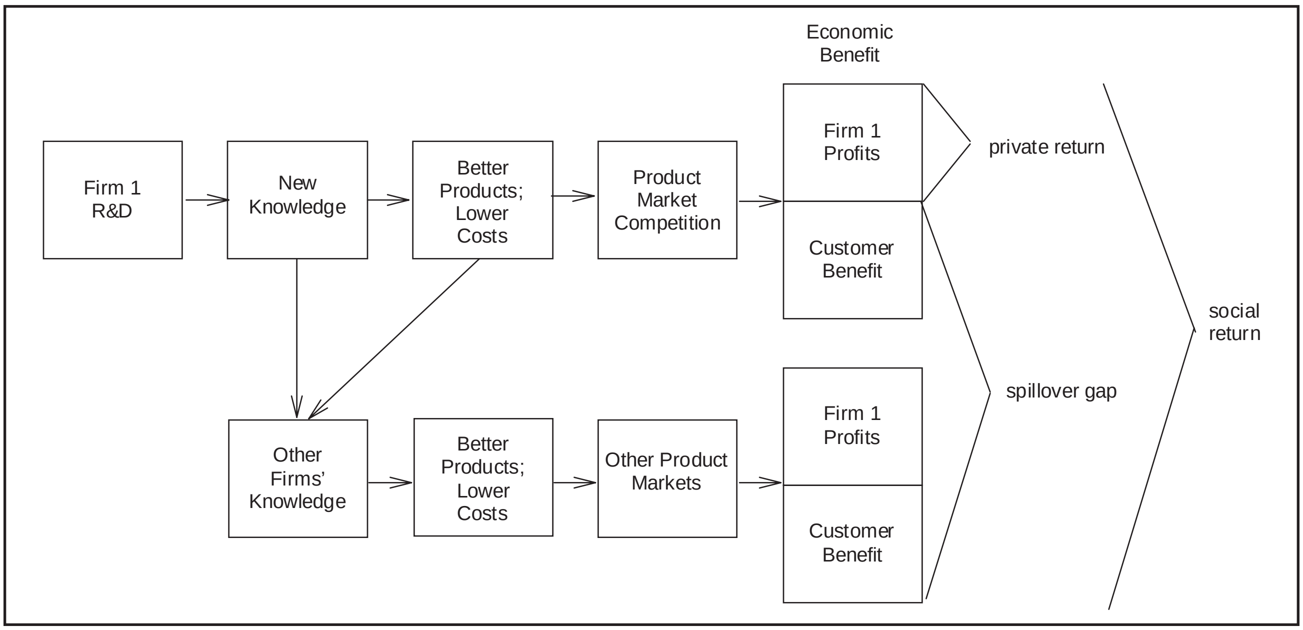

Figure 2. Private and social returns to R&D: Pure market spillover plus pure knowledge spillover

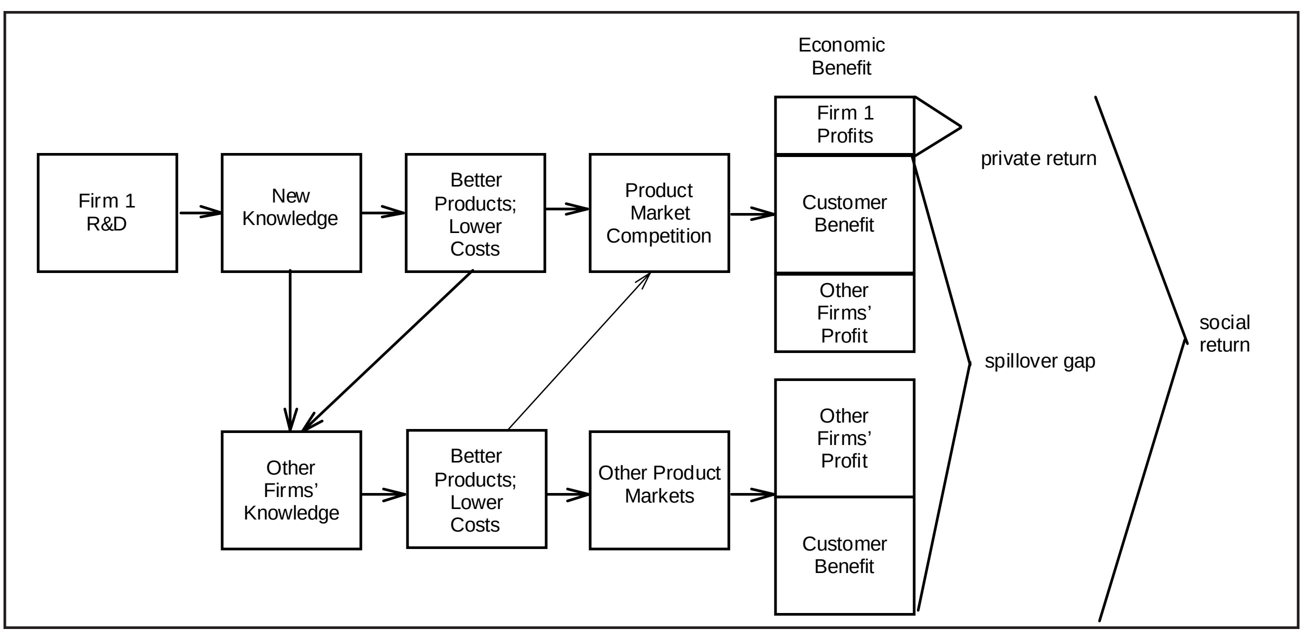

Figure 3. Private and social returns to R&D: Pure market spillover plus pure knowledge spillover plus interaction of the two

1" invests in R&D, this generates new knowledge, leading eventually to improved products or lower production costs. The operation of competition in the markets where firm 1's products are sold will divide the economic benefit of these improvements between firm 1's profits and benefits captured by customers in the form of lower costs or higher quality. In some cases, an innovating firm may not be in a position to utilize its new technology, but will need to license or sell the technology to another firm before the product or process can be implemented. In this case, imperfections in the licensing market will generally result in an additional spillover to the licensing firm.

The total social return to the innovation is comprised of the customer benefit plus the profits accruing to firm 1; the private return is only firm 1's profits, and hence there is a "spillover gap" consisting of the customer benefit. The more competitive are the markets in which firm 1's products are sold, the greater will be the share of the economic benefit that will be driven out of firm 1's profit and into the benefits captured by firm 1's customers. It is obvious from Figure 1 that the market spillovers will not be realized unless the innovation is commercialized successfully. Market spillovers accrue to the customers that use the innovative product; they will not come to pass if a technically successful effort does not lead to successful commercialization.

Note that market spillovers occur whether the purchaser of the new product is a household or another firm. In the case of improved intermediate products, then the market spillover benefits will be passed to the purchasing firms, which will in turn tend to pass at least some of this benefit to their customers. An important case of market spillovers associated with intermediate goods is where the innovation is an input to the research process, such as a new material or instrument. The purchaser is another researcher, who will typically use the new device in ways that create further spillovers.

Pure Knowledge Spillover. Figure 2 illustrates the effect of adding a "pure" knowledge spillover. By "pure," I mean a knowledge spillover that flows to firms that do not compete in firm 1's markets. Their increase in knowledge as a result of firm 1's research allows them to improve their products or lower their costs, increasing their profits and customer benefits in their markets. Both these profits and the consumer benefits are part of the social return, but are not captured by firm 1, and so the spillover gap is increased.

Note that even in the case of knowledge spillovers, the social return is created by the commercial use of a new process or product, and the profits and consumer benefits thereby created. The difference between market spillovers and knowledge spillovers is that in the former case the commercial benefits are created in the market for the new product or process that is the direct "output" of the research effort, while in the case of knowledge spillovers the commercial benefit is created indirectly through the creation of a new or improved product or process in some other market. Though as a society we value "knowledge for knowledge's sake," I am not including such non-economic value within the concept of knowledge spillovers used here.

Figure 2 indicates that the knowledge spillovers flow to some extent from firm 1's creation of new knowledge, and to some extent from firm 1's commercialization efforts. This reflects the idea that other firms may learn to some extent from papers, patents, departing employees, and other disembodied outputs of firm 1's research, but they are likely to learn more when firm 1's research results are actually embodied in new commercial products and processes. The relative importance and the speed of these two pathways will vary, depending on the nature of the research. In general, knowledge spillovers from more basic research would be expected to flow mostly from the research results themselves, and to take a fairly long time to have the commercial impact indicated in the lower part of Figure 2. On the other hand, knowledge spillovers from applied research and development are more likely to flow from the products or processes embodying the research results, and thereby have a quicker economic impact.

Thus, for the kinds of applied research and development projects that are the focus of the ATP, the realization of spillover benefits, and social returns more broadly, is strongly dependent on successful commercialization of the new technology. This is true both for market spillovers (which depend entirely on commercialization) and knowledge spillovers (which are likely to be largely dependent on commercialization). New products and processes that remain "on the shelf" do not benefit customers and hence do not create market spillovers, and their knowledge spillover impact is likely to be limited and/or distant in time. Basic research of the sort that is the mission of other federal agencies besides ATP is likely to create knowledge spillovers that are more diffuse and much more long-term.

The Interaction of Market and Knowledge Spillovers. It will often be the case that at least some of the firms that benefit from the knowledge spillover will be competitors or potential competitors of firm 1. The extreme case of this is pure imitation, where other firms copy the innovations of firm 1; more generally, firms making similar or related products may be able to improve their products or lower their costs on the basis of things they are able to learn as a result of firm 1's research. As shown in Figure 3, this complicates the picture in two ways. First, the introduction of these cheaper or better products into firm 1's markets creates some additional customer benefits, and some profits for these other firms, both of which constitute social returns not captured by firm 1. These increase the spillover gap.

At the same time that this increased competition increases social returns, it will likely reduce firm 1's profit from its own innovation. That is, the combination of knowledge spillover with competitive interaction increases the spillover gap both by raising the social return and lowering the private return. Thus "pure" knowledge spillovers increase the social rate of return to R&D, but they do so in a way that at least does not reduce the private return. When knowledge spillovers are combined with competition, however, the effect is likely to be an actual reduction of the private rate of return. Put differently, the interaction of knowledge spillovers and market spillovers aggravates the firm's appropriability problem: not only does the firm create benefits that it cannot capture, but its own profits from marketing its innovation are competed away. Understanding this interaction has important implications for identifying which research projects are likely to have large spillovers. In the section below, I discuss the factors that economists have identified that affect firms' ability to deal with this appropriability problem.

Framework for Explicit Evaluation of Spillover Potential of ATP Proposals

The Underlying Criterion for Project Selection

It is a generally accepted criterion of public policy that expenditure programs should seek to maximize the social rate of return of the expenditures they make. Maximizing the social return on ATP's investment is complicated by the possibility that ATP funding may partially or wholly displace private R&D resources, implying that the social benefits of the research would have come about without ATP. The possibility of displacement induces a distinction between the social rate of return to the project and the social rate of return to the ATP funds. If ATP funds a project with a high social rate of return, but in so doing largely displaces private funds, then the social return to the ATP expenditure will be low despite the high social return to the project.

Third-party surveys sponsored by the ATP, statistical analyses of the ATP's database of direct reports from participating companies, and a recent study by the General Accounting Office (GAO), all show ATP grantees believe that the great majority of ATP-funded projects would not have been undertaken, or the project schedules would have been slowed and project goals delayed without ATP funding. Note, however, that even if ATP funding accelerates the project, partial displacement could still be going on. Suppose, hypothetically, that the private proponents would have spent $500,000 per year, and the budget with ATP support is $600,000 with 50/50 cost sharing. In such a case, the project would be accelerated, but $300,000 in public funds would produce only a $100,000 increase in research effort. Each public dollar would, in this example, correspond to only 33 cents of increased project funding.

The danger of displacement means that what ATP must try to do is fund projects that have a high social rate of return, and a low probability that ATP funds are displacing private funds. Of course, ATP can never know for sure the extent to which it is displacing private funds, and project proponents have an inherent incentive to understate the likely extent of displacement.

Minimizing Displacement by Maximizing the Spillover Gap

The path through this dilemma is to look for factors that cause the social and private rates of return to diverge: the presence of such factors signals the possibility that social returns may be high at the same time that the risk of displacement is low. Strong likelihood of research spillovers is just such a factor. Hence by trying to identify project proposals where the likelihood of spillovers is particularly high, ATP will fulfill its statutory mandate, and do so in a way that will yield a high social return by minimizing the extent of displacement.

The relationships among the social rate of return, the private rate of return and the danger of displacement are illustrated by Figure 4, which graphs the social and private rates of return for various hypothetical projects. Obviously, there will always be tremendous uncertainty ex ante about the private and social returns to a project. Conceptually, Figure 4 should be thought of in terms of the expected returns, i.e., the magnitude of the return if successful, times the probability of success. The public sector seeks to maximize the expected social return, and the private sector seeks to maximize the expected private return.

Since projects higher up on the diagram have higher social returns, in the absence of the displacement concern and other constraints, ATP would simply seek to find projects that are as far up as possible. From the private sector point of view, projects to the right (higher private return) are more likely to be funded, all else equal. Of course, the likelihood of private funding for any particular project will depend on its riskiness and the financial environment of the project proponents. Although it is a gross oversimplification, for the purposes of discussion I have arbitrarily divided the projects into 3 groups: "good" commercial prospects that are likely to be well-supported by the private sector, "marginal" commercial prospects that are less likely to be funded and may be funded at inadequate rates, and "poor" commercial prospects.

All projects such as "A," "B," and "C" that lie above the 45° line generate spillovers. (Their social rates of return exceed their private rates.) If ATP seeks to choose projects with the highest social rate of return, then project "C" is the most desirable of these projects, and ought to be the prime candidate for funding. If society as a whole faced an all or nothing choice among these projects, we would indeed want to choose C, since its overall social rate of return is higher. But it is likely that C will be funded by the private sector, whereas it is likely that A and B will not be, or will be underfunded. If ATP ranks projects based on the "spillover gap," then projects A and B would indeed be favored over C. Hence if we want ATP to generate high returns from projects that would not otherwise be funded, then we would be better off looking at the spillover gap than the overall social return.

Project D illustrates the extreme version of this problem. This hypothetical project generates high social and

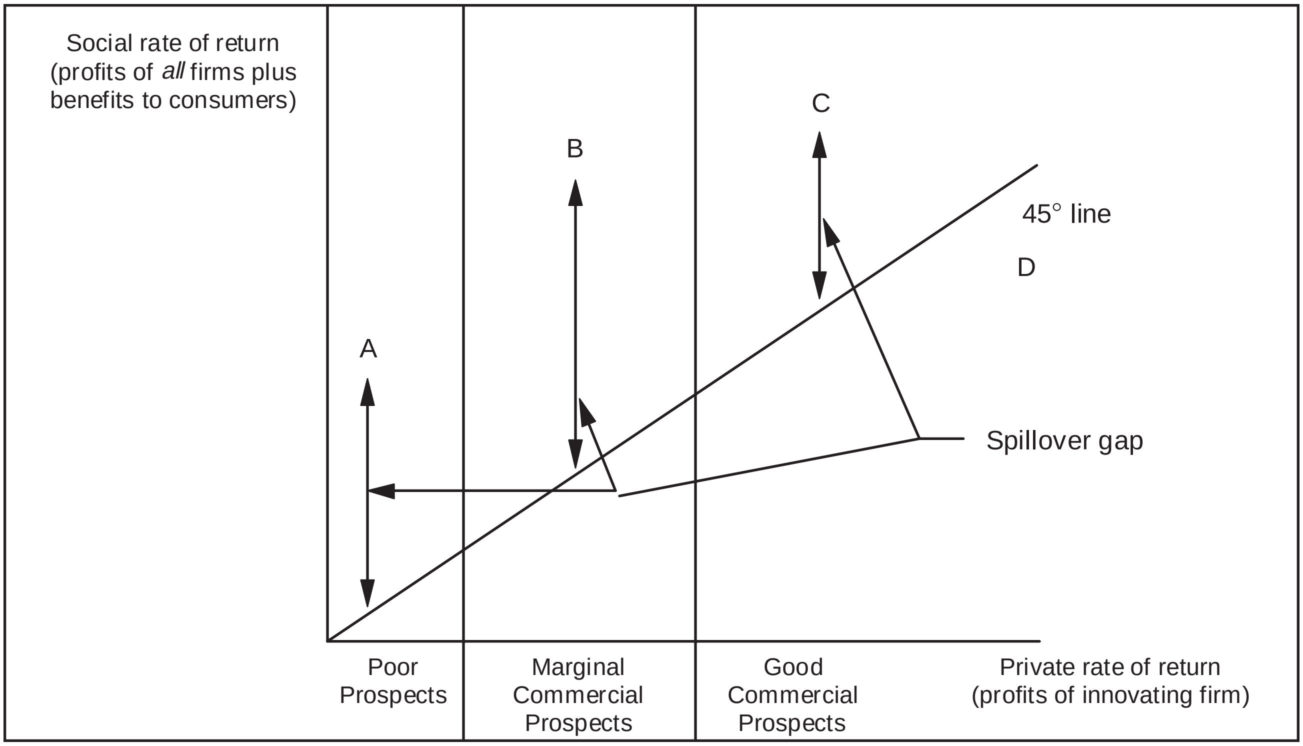

Figure 4.

private returns, but its net spillovers are negative. This might be the profile of a product that is highly successful but drives out a close substitute technology. If we were to seek to select projects only on the basis of social returns, this project would rank as highly as A, despite the fact that its private rate of return exceeds its social rate. The government has no business funding this project, however, given that its net effect on all parties, other than the firm performing it, is negative.

If ATP ranks projects based on the "spillover gap," than projects A and B are both attractive. But project A, which has large social returns but very low private returns, may not be a viable ATP project in practice. If submitted, it would score high on ATP criteria that emphasize potential spillover benefits, but the ATP would face the dilemma of a weak business case and an uncertain path to commercialization and the achievement of the potential benefits. Hence projects such as B, which have a significant spillover gap but also offer some prospect of private returns, are the most appropriate candidates for the ATP.

In reality, of course, we will have only coarse estimates of the social rate of return or the spillover gap. The fact that these prospects can only be known with great uncertainty strengthens the superiority of the spillover gap over the social rate of return as a decision criterion. Although projects like A do exist, there will, in general, be some correlation between the private and social rates of return. For example, all else equal, both rates will be higher for projects with higher success probabilities and projects whose product (if successful) serves a larger market. If we focus only on the social rate of return, then there is a danger that we will fund projects that appear to have a high social rate of return, where the only reason the private sector is not pursuing this project is because its overall prospects (affecting both the private and social returns) are being overestimated or overstated by the project proponents. If we focus on the spillover gap rather than the apparent overall social rate of return, we are less likely to step in to fund projects for which the explanation for lack of private funding is that they are not really very promising projects.

To state this point slightly differently, the ATP decision process should recognize that its information is imperfect, and that errors are going to be made. Further, information that is available about a project ought to be examined not only for what it says explicitly about social returns, but also for it what it implies about the probability of errors being made. In some cases, the "facts" being put forth to support the likelihood of large social returns for a project are facts that equally well support the likelihood of large private returns (e.g., a large market for the resulting product). If these purported facts are true, then both private and social returns from the project will be high, i.e. we will be at points like C (or even D) in Figure 4. Now, the ATP cannot know with certainty if the "facts" are really true, and cannot know why, if they are true, the private sector would not fund the project on its own. Logically, there are several possibilities: (1) the facts as presented are actually false, i.e., the market is not really large or else the probability of success is so low that the expected (social) net present value of the investment is negative; (2) the facts as presented are true, and the private sector knows it and would, indeed, fund the project with or without government help; or (3) the facts are true, and there is some reason why the private sector will not fund the project (or will not fund it adequately or in a timely way) despite the potential payoff.

Some projects will, of course, fall into category (3), but the ATP should be worried about possibilities (1) and (2). This worry can be minimized by seeking a large spillover gap, not just a large social return. If projects with apparently high potential private returns are to be funded, there should be a careful analysis of the reasons as to why the project is not being funded despite its large potential payoff.

Measuring Rates of Return and Spillovers in Impact Studies

A possible excuse for the delay between the time Alfred Marshall talked about spillovers and the time economists made serious efforts to measure them is that they are inherently difficult to observe. As Paul Krugman has noted, "knowledge flows . . . are invisible; they leave no paper trail by which they may be measured and tracked, and there is nothing to prevent the theorist from assuming anything about them that she likes" (Krugman 1991, p. 53). As a result, empirical measurement of spillovers is necessarily somewhat indirect. Most analyses take the form of measuring the innovative effort or output of one agent or set of agents, and looking for a correlation between this measure and the innovative output of another agent or set of agents.

To make such an analysis tractable and meaningful, one must identify which agents are the likely recipients of spillovers from particular research efforts. This typically involves developing a metric for measuring the "closeness" of different agents—either in terms of technological similarity, geographic proximity, or economic relationships, such as vendors and their customers. To infer the existence of spillovers from a correlation between the research effort of one group of agents and the research output of other agents that are somehow "close," it is necessary to control for (1) the innovative effort of the second group, and (2) variations in "technological opportunity" that might be affecting the productivity of research effort for both the "spilling" and "receiving" agents, inducing a correlation between an agent's research success and the effort of other firms that need not be related to spillovers.

Studies of this sort allow the calculation of the "excess return" to R&D investment, i.e., the difference between the rate of return calculated including the effects of the investment on the recipients of spillovers, and the rate of return calculated excluding spillover effects. Depending on the nature of the study, this excess return or spillover gap may encompass knowledge spillovers, market spillovers, or both. In general, the spillover gap is found to be positive, suggesting that the negative competitive externality is generally outweighed by positive effects of knowledge and consumer surplus externalities.

Measurement of Market Spillover

The oldest line of work focuses on spillovers embodied in products and measures closeness using supplier-customer relationships. For example, Terleckyj (1974) looked at industry data, constructing a measure of "borrowed" R&D for each industry on the basis of the R&D of the industries from which it purchased intermediate inputs, including capital equipment. He found that the productivity effects of R&D in downstream industries implied an excess return to industry R&D of 20% to 50% (compared to a private rate of return of about 30%). This measure of market spillovers may also contain an element of knowledge externalities, to the extent that the downstream firms are engaged in their own research and benefit indirectly from the research of their suppliers.

Scherer (1982 and 1984) took another cut at this problem. By examining patent data, he estimated the fraction of inventions originating in each industry that would be used by each industry. This allows the creation of a "technology flow" matrix which can be used to allocate industrial research by the industry in which it will be "used" regardless of the industry in which it is performed. He shows that this "used" R&D variable is more strongly correlated with industry productivity growth than is a variable measuring R&D performed in the industry.

Mansfield, et al. (1977) used a case study approach instead of looking at aggregate industry R&D statistics. They identified 17 specific innovations, and attempted to estimate the actual cost and overall social benefits of each. In particular, they took great care to analyze the impact of the innovations on customers, and also on competitors. They did not, however, specifically seek to identify knowledge externalities. For this group of innovations, the median private rate of return was about 25% and the median social rate of return was about 50%.

Bresnahan (1986) and Trajtenberg (1990) have quantified the consumer surplus spillover from mainframe computers in the 1960s and the CT scanner in the 1970s. Bresnahan calculates that between 1958 and 1972 financial service firms paid $68 million for computing services, but received benefits equal to $200 to $400 million. Although he does not explicitly calculate rates of return, this clearly shows that the social rate was several times the private rate. Trajtenberg calculates that the social rate of return to improvements in CT scanners averaged between 180 and 350 percent per year, depending on how foreign R&D investments are treated. While Trajtenberg also does not calculate private rates of return, approximately half of the producers, including EMI, the original innovator, eventually left the business, apparently because of mounting losses.

Measurement of Knowledge Spillovers

In my 1986 paper, I used patent data for about 500 manufacturing firms to characterize the "technological proximity" of all pairs of firms on the basis of the extent of overlap of technological classification of their patents. I then constructed a measure of the "spillover pool" for each firm, as the sum of all other firms' R&D, weighted by their proximity to the receiving firm. I found that the pool variable had positive effects on firms' patents, profits and market value, all controlling for the firm's own R&D. For patents—a purely technological measure of research output—roughly half of the aggregate impact of R&D was in the form of spillovers, or, conversely, the social productivity of research was roughly twice as great as the private productivity. For economic measures of research output such as profits, productivity and market value, I found that the spillover effect was roughly half as large as the private return.

Interestingly, the effect of the pool was found to be itself a function of firms' own R&D. The more R&D a firm does itself, the more it benefits from spillovers from others. With respect to profits and market value, firms that have significantly less than the mean R&D level actually suffer a negative effect from the spillover pool. This is interpreted as saying that both knowledge and competitive externalities are present, with the former outweighing the latter on average, but the latter outweighing the former for firms that do little R&D themselves.

Summary of Estimates of Spillover Magnitudes

Griliches (1992, Table 1) summarizes the results of many of these studies. He concludes "R&D spillovers are present, their magnitude may be quite large, and social rates of return remain significantly above private rates." While all of this work carries econometric limitations and presents only indirect evidence that spillovers exist, the weight of the evidence does seem to be increasingly convincing that spillovers create a large gap between the private and social rate of return. There are two ways to look at this gap. In absolute terms, it appears that the excess of the social rate of return over the private rate—the rate of spillover—is something like 15 to 30 percent, with some estimates much higher than that. Another way to look at this is relative to the private rate of return. Again, estimates vary somewhat, but spillovers seem to create a gap between the private and social return that is equal to 50 to 100% of the private rate of return. Note that the individual studies underlying these ranges tend to emphasize either knowledge externalities or market externalities. I can think of no study that, at a conceptual level, is designed to capture both, although relationships between the two in the data make it likely that each kind of study picks up some of the other effect. Hence it is likely that these estimates have some tendency to underestimate the combined effects.

Conclusion

In order to be effective in achieving its statutory objectives, ATP must try to determine which projects proposed to it will generate large spillovers, and which will not. Economists and other social scientists have identified certain aspects of a project's technological and market environment that tend to be associated with large or small spillovers. By incorporating the explicit analysis of such factors into both project choice and evaluation of project impact, ATP can make better decisions

This is an inherently difficult and uncertain task, and it is one that requires an unusual combination of technical, business and economic analysis. Perfect prediction cannot be achieved, any more than it can be achieved for the purely technical success of research. We know enough about spillover prediction and measurement to improve ATP's project selection and evaluation of outcomes using more systematic, explicit treatment of spillover effects. Further research can improve and extend our knowledge of spillover phenomena and how to measure them, in order to provide a firmer foundation for a program with the mission, goals, and strategies of the ATP.

The empirical evidence suggests that the average research project generates spillovers. If ATP can succeed in targeting projects with better-than-average spillover potential, then it will generate large social returns that would not otherwise have been achieved.

References

Bresnahan, T. "Measuring the Spillovers from Technical Advance: Mainframe Computers in Financial Services." *American Economic Review,* September 1986.

Cohen, Wesley, and Daniel Levinthal. "Innovation and Learning: The Two Faces of R&D." *Economic Journal,* 1989.

General Accounting Office. *Performance Measurement: Efforts to Evaluate the Advanced Technology Program.* GAO.RCED-95-68, 1995.

General Accounting Office. *Measuring Performance: The Advanced Technology Program and Private-Sector Funding.* GAO/RCED-96-47, 1996.

Griliches, Zvi. "The Search for R&D Spillovers." *Scandinavian Journal of Economics,* 1993.

Jaffe, Adam B. "Technological Opportunity and Spillovers of R&D." *American Economic Review* 76, December 1986, pp. 984–1001.

Krugman, Paul. *Geography and Trade.* Cambridge: M.I.T. Press, 1991.

Levin, Richard C., Alvin Klevorick, Richard Nelson, and Sidney Winter. "Appropriating the Returns from Industrial Research and Development." *Brookings Papers on Economic Activity,* 1987, pp. 783–831.

Mansfield, Edwin, J. Rapoport, A. Romeo, S. Wagner, and G. Beardsley. “Social and Private Rates of Return From Industrial Innovations.” *Quarterly Journal of Economics*, 1977.

Marshall, Alfred. *Principles of Economics.* London: Macmillan, 1920.

Trajtenberg, Manuel. *Economic Analysis of Product Innovation: The Case of CT Scanners.* Harvard University Press, 1990.

Author Biography

Adam Jaffe is a Professor of Economics at Brandeis University. His research interests include the effect of incentives and regulations on the diffusion of technologies, and the role of R&D and innovation in economic growth. Jaffe coordinates the National Bureau of Economic Research project on U.S. productivity and industrial technologies.